

# Overview of Implementing ATM-Based Enterprise Local Area Network for Desktop Multimedia Computing

Yiwei Thomas Hou, Polytechnic University  
Leandros Tassioulas, University of Maryland  
H. Jonathan Chao, Polytechnic University

**A**synchronous transfer mode (ATM) offers much greater capacity than existing shared-medium local area networks (LANs). It is a flexible transport technology that can integrate all types of media such as voice, video, and data. Many of the key features of ATM have been standardized by the ATM Forum for the purpose of local area desktop computing. Various LAN vendors are actively developing and implementing ATM products in the market.

In this article, we present a unified overview of the ATM-based enterprise network (or ATM LAN). We will cover key components and issues regarding the implementation of an ATM-based LAN and use commercial ATM LAN vendors' products to illustrate. Also, we will cover the latest ATM Forum standardization efforts on traffic management functions and LAN emulation. This article is organized as follows. We first outline the major features of an ATM switch that should be taken into consideration when evaluating the ATM switch product used as an enterprise network backbone. Next, we examine the strategies of interconnecting host computers with an ATM switch (i.e., introducing ATM onto the desktop). Then we present the ATM traffic service classes and the associated traffic management functions. Finally, we discuss how to seamlessly support the existing Transport Control Protocol (TCP)/Internet Protocol (IP) in an ATM environment.

## PERFORMANCE FEATURES AND CHARACTERISTICS OF ATM SWITCHES

Intensive research on ATM switching architectures has been done for the past decade [1]. Some of the architectures proposed in the literature have been developed by vendors and are available in the market now. For a corporate network planner, it is essential to evaluate these commercial ATM switches within some guidelines. In this section, we present major factors that should be considered when an ATM switch

---

*The authors discuss important implementation issues in an ATM-based enterprise network, and propose possible migration strategies for the smooth introduction of ATM into the desktop computing environment.*

---

is evaluated. These include alternative switching architectures, storage buffer management, multicasting, scalability, traffic performance guarantees, and so on [2].

### TIME AND/OR SPACE DIVISION SWITCH

An ATM switch can be classified as time division, space division, or a combination of both (Fig. 1 [3]).

In a switch fabric based on time division, all cells flow across a single communication highway shared by

all input and output ports. This communication highway may be either a shared medium or a shared memory. The throughput of this single shared medium or memory defines the capacity of the entire switch fabric, and fixes an upper limit on the capacity for a particular implementation beyond which that switch cannot grow. Since every cell flows across a single shared resource, this class of switch may easily support multicast operation.

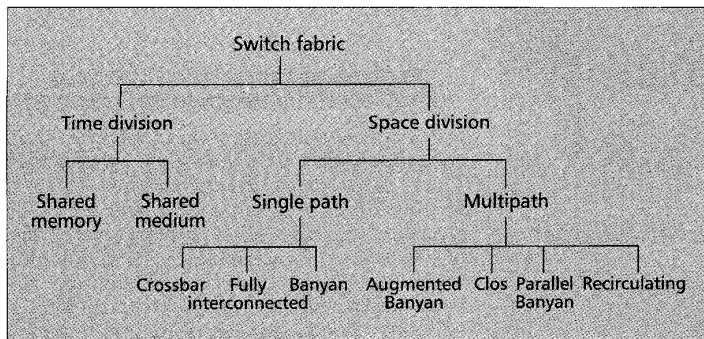
Figure 2 illustrates the functional block diagram of a 16 x 16 time division multiplexing (TDM) shared-medium ATM switch. Among commercial ATM LAN vendors, switches based on TDM shared medium are Adaptive Corp.'s ATM switch [4], FORE Systems' ASX-200 switch, Newbridge Networks' VIVID switch, and others. In space division, a multiplicity of paths are provided between the input and output ports. These paths operate concurrently so that many cells may be transmitted across the switch fabric at the same time. The total capacity of the switch fabric is thus the product of the bandwidth of each path and the number of paths that can transmit a cell concurrently. The upper limit on the total capacity of the switch is therefore theoretically unlimited. In practice, however, it is restricted by physical implementation constraints such as device pin count, connector restrictions, and synchronization considerations, which together limit the size of the switch fabric. Commercial ATM LAN switches based on space division are the HSS switch from Alcatel Data Networks, LattisCell switch (Fig. 3) from Bay Networks, and others. There are trade-offs between time-division- and space-division-based switch designs in terms of bus and memory speed requirements, scalability of switch size, internal link blocking, and other aspects.

### MEMORY SPEED REQUIREMENT

Since memory access time is constrained by the hardware technology, this will set a limit on ATM switch size. Switches

---

*This work is supported by a National Science Foundation (NSF) Graduate Research Traineeship, in part by NSF grant NCR-9406415 and AFOSR contract F49620-95-1-0061, and by the New York State Center for Advanced Technology in Telecommunications at Polytechnic University, Brooklyn, New York.*



■ Figure 1. Classification of ATM switching architecture.

based on TDM require memory access speed on the order of  $N \times$  port bandwidth. Therefore, their switch size will be the most affected by the memory speed constraint. For example, for a TDM shared-memory ATM switch, let the clock frequency be 100 MHz, or

$$t_c = \frac{1}{100 \text{ MHz}} = 10 \text{ ns/cycle}$$

for each memory access, then we have the inequality

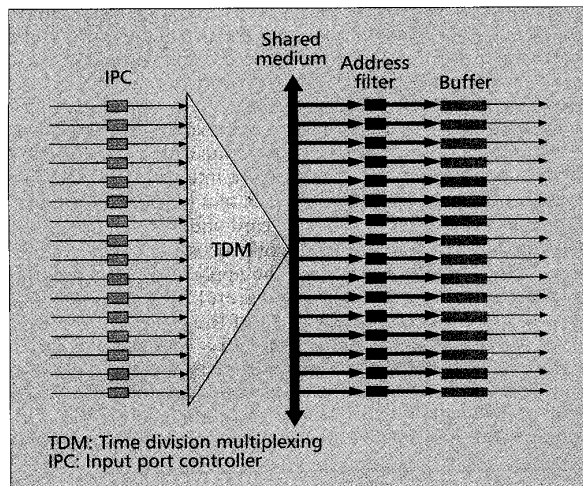
$$\frac{2(NB)}{P} \leq \frac{1}{t_c}$$

where  $N$  is the switch size,  $B$  is the port bandwidth,  $P$  is the bus width, and the 2 reflects memory access for both read and write operations. If  $P = 128$  and  $B = 160 \text{ Mb/s}$  (including routing overhead for OC-3), then  $N \leq 40$  ports. That is, the switch size is at most 40 if it is based on TDM shared memory.

For a space-division switch, each switch module operates on the same order of each port bandwidth and is therefore relatively unaffected by the memory speed constraint.

#### INTERNAL LINK BLOCKING AND OUTPUT PORT BUFFER CONTENTION

Internal link blocking occurs when multiple cells contend for a link at the same time *inside* the switch fabric. This usually happens in a switch based on space division when an internal physical link is shared by multiple connections among input/output ports. For example, the LattisCell switch (Fig. 3) from Bay Networks operates at 155 Mb/s on each set of



■ Figure 2. Time division multiplexing shared-medium ATM switch.

input/output ports. All internal switching fabric links run at 310 Mb/s (speedup factor of 2) to help reduce blocking. However, blocking can still occur because at a later stage of switching fabric, it is possible that cells from two links, each carrying 310 Mb/s traffic, compete for the same output link. To reduce cell loss, buffering is used at each binary switch element throughout the fabric.

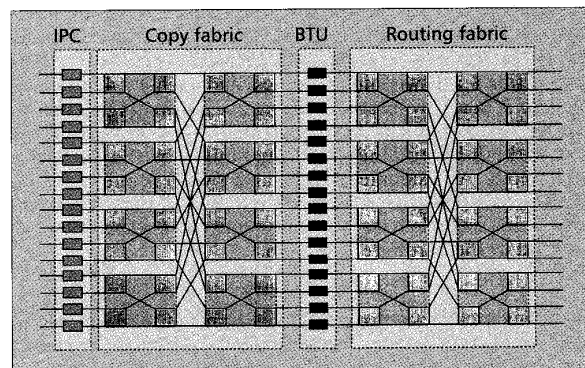
Output port contention occurs when a number of cells arrive simultaneously from different input ports, each requesting the same output port. A single output port can only transmit one cell at a time; thus, only one cell can be transmitted, and the others routed to that output port must be either discarded or buffered. Output port contention resolution is present for every design of ATM switches.

#### CELL ROUTING MECHANISM (SELF-ROUTING OR LABEL ROUTING)

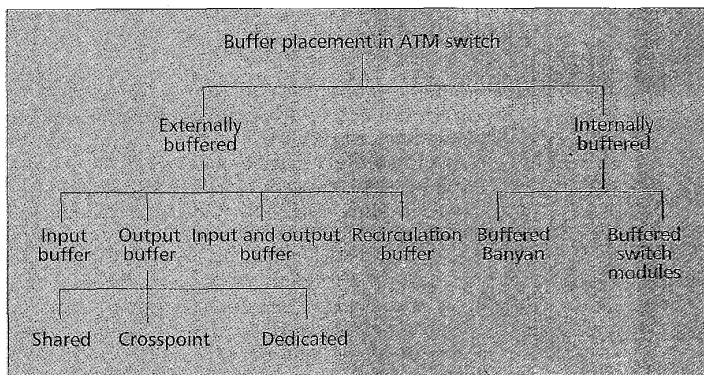
Self-routing uses a special binary header (in addition to the five-byte cell header) that is prepended to each cell at the input port before it enters the switching fabric. The extra binary header contains the information used to navigate cells through the switch fabric to their destined output ports. Binary or self-routing works by sequentially examining each bit of the special header and switch the cell to either the upper output of the switching element if the bit is 0 or the lower output of the switching element if the bit is 1 (e.g., Banyan family switches such as Banyan, Delta, and Omega). The special binary header is, of course, removed at the output port after the cell traverses the switch fabric and is therefore transparent to user traffic. In contrast, in label routing the virtual channel identifier (VCI) field within the header is used by each switching element to make the output link decisions. That is, each switch module stores a VCI in a lookup table and switches cells to an output link according to the mapping between the VCI and input/output links in the table. Label routing does not depend on the regular interconnection of switching elements the way self-routing does, and can be used where switching elements are interconnected arbitrarily.

#### SINGLE-PATH OR MULTIPATH

Single-path provides a unique internal link path for any given input/output pair. Examples are Banyan family switches such as Banyan, Delta, and Omega, which are characterized by  $\log_2 N$  switching stages ( $N$  is the switch size). Internal blocking is a problem with these single-path switches. To alleviate this problem, a multipath switch provides multiple alternative internal switching paths for a given pair of input/output ports.



■ Figure 3. Bay Networks' LattisCell ATM switch architecture.



■ Figure 4. Buffering strategies for an ATM switch.

A well known example is the Benes network, which is based on Banyan but has added stages to provide multiple alternative paths. A multipath switch reduces internal link blocking with a connection reconfiguration algorithm. However, it has the potential of causing out-of-sequence cells due to differences in delay among alternative paths.

#### BUFFERING STRATEGY (BUFFER PLACEMENT) AND UTILIZATION EFFICIENCY (MEMORY SHARING)

Buffering is required for ATM switch design when multiple cells compete for the same output link at a switch element or output port. The location of the buffers has a major effect on the overall performance of the switch and also affects the complexity of the switch fabric. Figure 4 [3] gives a classification of buffering strategies for ATM switches. The more sharing of memory/buffer among the ATM traffic, the less buffer space is required; but on the other hand, the complexity of control logic goes up as we get more out of the buffer sharing.

#### HANDLING OF MULTICAST CONNECTIONS

Multicast/broadcast is essential for one-to-many communication (e.g., desktop video conferencing). Some switching fabrics achieve multicast by first replicating multiple copies of the same ATM cells and then routing each copy to its intended output port (e.g., the LattisCell switch from Bay Networks [5]). Other switches achieve multicast by utilizing the inherent broadcasting nature of shared-medium without generating any copies of ATM cells (e.g., FORE's ASX-200 switch [6]).

#### MODULARITY AND SCALABILITY

In order to implement ATM switches of moderate to large size, a modular approach is necessary in which a hierarchy of switching components is created. The most fundamental switching elements are interconnected to form a larger switching structure called the *switch module*. Multiple switch modules can be connected to form an ATM switch. For example, the LattisCell 16 x 16 ATM switch (Fig. 3) is made by interconnecting 4 x 4 switching modules, and each 4 x 4 switching module is made by interconnecting 2 x 2 switching elements. An ATM switch is said to be scalable if it can grow to larger size without being fundamentally limited by the hardware technology. An example of a scalable switch is one based on space division (Fig. 3). A switch based on TDM (shared-medium) is not scalable because the memory access speed is constrained by the hardware limitation.

#### COMPLEXITY

Switching architecture complexity in terms of total number of interconnection stages, total number of switching elements, and, most important, the number of interconnection wires is a measure of implementation cost.

#### EASINESS OF CLOCK SYNCHRONIZATION (UNIFORM INTERCONNECTION WIRE LENGTH?)

It is a desirable feature to have identical short connection wires between switch elements [7]. This is because timing alignment for signals at identical short connection wires between switch elements (e.g., crossbar) is much easier than that of other types of interconnection (e.g., Shuffle). The unequal length of interconnection wires increases the difficulty of synchronizing the signals and consequently limits the switch fabric's size (e.g., Batcher-Banyan network).<sup>1</sup>

#### THROUGHPUT/SWITCHING DELAY AND DELAY JITTER/CELL LOSS PERFORMANCE

Different switching architectures provide different throughput performance. For example, for input port buffered switch, the throughput is at most 58.6 percent for uniform uncorrelated traffic due to head of line (HOL) blocking, while for output buffered switches the throughput can reach 100 percent.

The time to switch an ATM cell through the switch is called the *switch transit delay*. Typical values mentioned in the literature range between 10 and 1000  $\mu$ s, with a jitter of a few hundreds of microseconds or less [8].

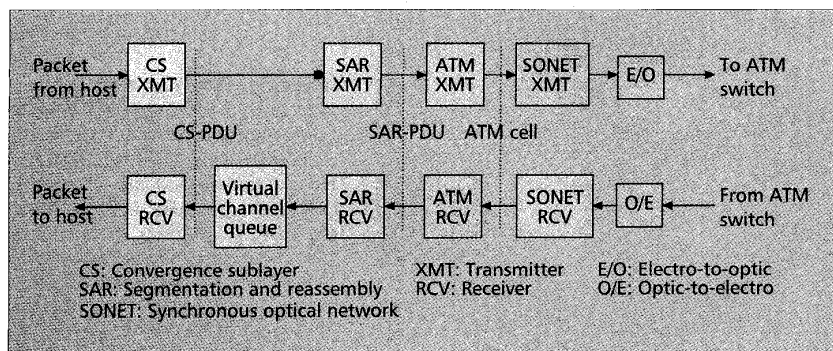
Cell loss performance is critical for data applications. For connections using the TCP/IP protocol, once an ATM cell (segmented from an IP packet) is lost, the whole IP packet must be retransmitted if a reliable transport protocol such as TCP is used, which may introduce excessive overhead traffic to the switch. Typical requirements for cell loss probability in ATM switches is between  $10^{-8}$  and  $10^{-11}$  [8].

#### INTERCONNECTING DESKTOP COMPUTERS WITH AN ATM SWITCH: THE HOST/NETWORK INTERFACE

One way to introduce ATM into the desktop computing environment is to use an ATM adapter card on the host computer and connect the adapter card directly onto an ATM switch port. This will provide whole switch port bandwidth (e.g., 155Mb/s) to a desktop computer. Figure 5 illustrates the functional block diagram of an ATM adapter card [9]. Since multiple packets of different VCs may arrive concurrently in an interleaved fashion, these partially reassembled packets are temporarily stored in the virtual channel queue (VCQ). As soon as a packet is completely received, it is immediately sent to the convergence sublayer (CS) receiver and then to the host.

There are two important issues concerning the design of the host-workstation-to-ATM-network interface [10]. The first concerns the partitioning of the protocol (e.g., TCP) processing function between the adapter card and the host processor. Moving part or all of the protocol processing to the adapter gives better data throughput on the application level, but this is costly and requires extremely careful programming. The second concerns the integration of the adapter card with other host subsystems such as memory. The use of a direct memory access (DMA) engine will assist the movement of data between the host memory and the adapter card's buffer without involving the host central processing unit (CPU). This reduces the number of memory bus crossings.<sup>2</sup> Another

<sup>1</sup> For switches designed for LAN, interconnection wire length may not pose a problem since the switch size is small, typically under 100 x 100.



■ Figure 5. Functional block diagram of an ATM adapter card.

advantage of using DMA is that DMA transfers data to the host memory in parallel with CPU computation without interruption, which means fewer CPU cycles are lost.

An example of a commercial ATM adapter card is FORE's 200-series adapter (Fig. 6 [12]). As shown in Fig. 6, the network interface section includes the physical layer, which bridges the physical-medium-dependent layer, such as fiber, from the remainder of the adapter. The control processor section implements the ATM layer and the ATM adaptation layer (AAL). An Intel 960 processor performs AAL5 segmentation and reassembly, and manages the transfer of data between the adapter and the host computer. The bus interface section contains custom bus-specific hardware for the I/O bus of the host computer (e.g., SBus for Sun SPARCstations, GIO bus for Silicon Graphics workstations, etc.). A DMA engine is used to connect the adapter directly with the host memory I/O bus.

While it is necessary to use ATM adapter cards to connect powerful server machines and high-end workstations directly to the ATM switch to tap the high port bandwidth for better throughput, this may not be the case for most average desktop computers. The reason is twofold. First, unlike powerful server machines, most desktop computers' software throughput is likely to become a bottleneck when they are directly connected to the switch with a high-speed link, say OC-3c fiber. This is because the common implementation for a reliable transport protocol such as TCP/IP, which includes packetization, error handling, end-to-end flow control, routing, and congestion control, can achieve throughput of only a few tens of megabits per second [11, 13]. With this throughput, 155 Mb/s will certainly overwhelm the host CPU. Second, installing an ATM adapter card on the host computer requires replacing the existing Ethernet card on the host computer and possibly recabling the entire network. The cost of hardware and the rewiring workload will surely be too high to be feasible. Therefore, it will be essential to have a smoother

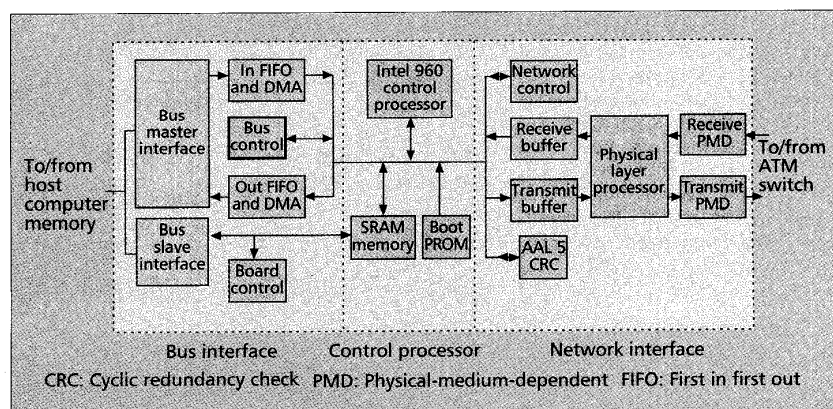
<sup>2</sup> A conventional network interface such as an Ethernet card will result in the data crossing the memory bus as many as six times for TCP processing for each data transfer [11].

<sup>3</sup> The use of an ATM hub may be analogous to the architectural concept of fiber to the curb (or rather to the wiring closet here) while the use of ATM adapter card is analogous to fiber to home (or desktop).

average user.<sup>3</sup> Rewiring is also minimal because only the part of the ATM hub to the switch is rewired with a medium such as fiber. This recabling operation will not be as physically noticeable to the users as the ATM adapter card.

A commercial example of ATM hub is Bay Network's EtherCell Ethernet-to-ATM switch, illustrated in Fig. 7 [14]. As shown in the figure, 12 10Base-T Ethernet ports are available to connect to host computers with 10 Mb/s dedicated bandwidth. To fully utilize the switching capability of ATM and reduce the cost of the ATM hub, traffic switching among ports on the same EtherCell ATM hub are performed by the ATM switch rather than by the hub itself. The ATM hub can be placed in the wiring closet which used to host the Ethernet hub. It is connected with fiber directly to the ATM switch in the central machine room. EtherCell uses proprietary CellFrame technology to perform hardware-based conversion between Ethernet frames and the ATM cell format. That is, Ethernet frames from the host computer are chopped up and formatted into ATM cells and switched through an ATM switch, while from the ATM switch cells get assembled into Ethernet frames in the EtherCell and delivered to the host computer on Ethernet.

The protocol stack conversion between a host and an ATM switch connected via an ATM hub or with an ATM adapter card is illustrated in Fig. 8, where LAN emulation is used (more on LAN emulation later). As shown in Fig. 8, the protocol conversion between an IEEE 802.x medium access control (MAC) frame and ATM cells is performed inside the ATM hub. When an ATM adapter card is used, the host is directly connected to the switch, and the AAL segmentation and reassembly is performed on the ATM adapter card.



■ Figure 6. FORE Systems' 200-series ATM adapter card architecture.

## TRAFFIC MANAGEMENT FUNCTIONS

The ATM Forum has classified the ATM layer service classes as follows [15].

### CONSTANT BIT RATE

Constant bit rate (CBR) is intended to support real-time applications requiring tightly constrained delay and delay variation. Example applications are voice, circuit emulation ( $N \times 64$  kb/s, T1, T3, etc.). The traffic descriptor used at call setup for CBR is peak cell rate (PCR). Quality of service (QoS) parameters specified for CBR are maximum cell transfer delay (CTD), peak-to-peak cell delay variation (CDV), and cell loss ratio (CLR).

### VARIABLE BIT RATE

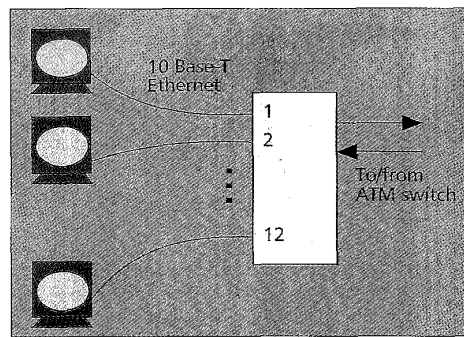
There are two subclasses of this service, namely real-time (RT) variable bit rate (VBR) and non-real-time (NRT) VBR. RT VBR is for applications requiring tightly constrained delay and delay variation such as video. Sources are expected to transmit at a rate varying with time. On the other hand, the NRT VBR service class guarantees a mean cell transfer delay and a cell loss ratio for those connections that remain within the traffic contract agreed on with the network at the time the connection is established. An example of NRT VBR application is response-time-critical transaction processing (e.g., airline reservations, banking transactions). Traffic descriptors for both RT VBR and NRT VBR are PCR, sustainable cell rate (SCR), and maximum burst size (MBS). QoS parameters for RT VBR are maximum CTD, peak-to-peak CDV, and CLR. QoS parameters for NRT VBR are mean CTD and CLR.

### AVAILABLE BIT RATE

Many data communications can tolerate reduction of their information transfer rate if the network requires it. Likewise, they may wish to increase their information transfer rate if there is extra bandwidth available within the network. Therefore, a service which utilizes the available network bandwidth (i.e., available bit rate — ABR) would be best for this type of service. Examples requiring ABR are data communications applications requiring a reasonable delay performance, such as remote procedure call, distributed file service (e.g., network file system — NFS), or computer process swap/paging. Traffic descriptors for ABR are PCR and minimum cell rate (MCR). The ATM network will dynamically adjust bandwidth between PCR and MCR depending on available network resources. QoS parameter for ABR is the CLR.

### UNSPECIFIED BIT RATE

A less reliable service than ABR, unspecified bit rate (UBR) is also called "best effort service." The UBR service is intended for delay-



■ Figure 7. Bay Network's EtherCell Ethernet-to-ATM switch.

tolerant or non-real-time applications (i.e., those that do not require tightly constrained delay and delay variations, e.g., file transfer submitted as a background job). The traffic descriptor for UBR is PCR. Since the UBR service class offers no traffic-related service guarantees, no QoS parameters are specified.

To successfully support the above ATM layer service classes according to their traffic descriptors and QoS requirements, certain traffic control functions are employed in ATM network. They

are *connection admission control, traffic shaping, traffic policing, and congestion control.*

### CONNECTION ADMISSION CONTROL

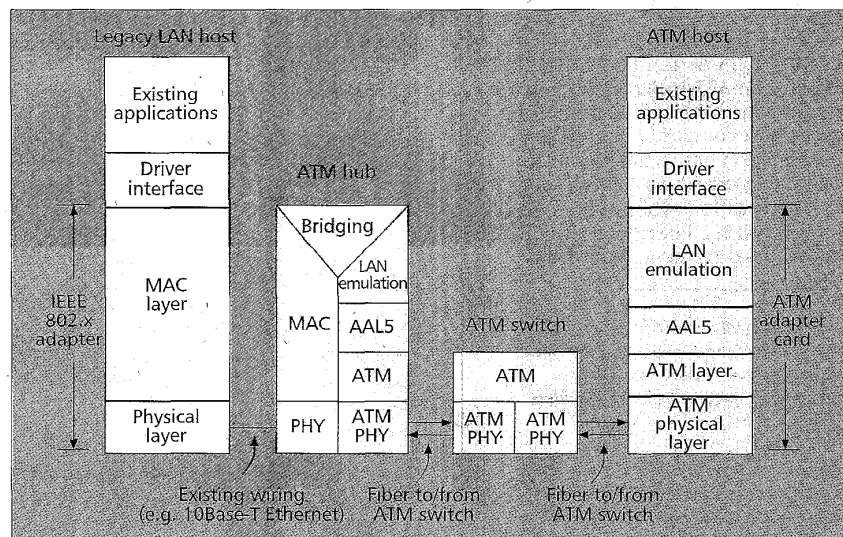
Connection admission control refers to the set of actions taken by the network at the call setup phase (or during the call renegotiation phase) in order to decide whether to establish a virtual channel/path (VC/VP) connection or reject the call request. A connection request is accepted only when sufficient resources are available to establish the connection through the whole network at the required QoS and to maintain the agreed QoS of existing connections.

### TRAFFIC SHAPING

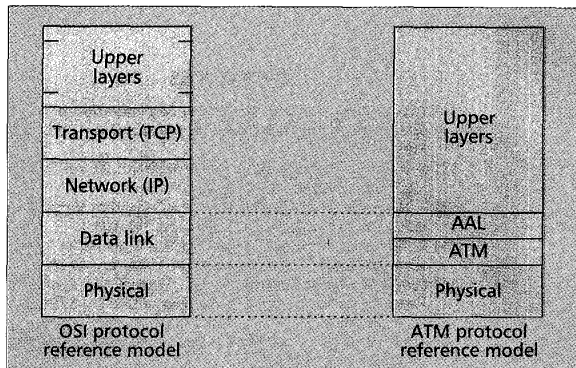
This is performed at the user-network interface (UNI) of the ATM network (e.g., ATM adapter cards and hubs) to ensure that traffic matches the contract negotiated between the user and the network. Traffic shaping ensures that traffic meets the parameters negotiated at connection establishment. Traffic is shaped according to the Generic Cell Rate Algorithm (GCRA) as specified by the ATM Forum UNI 3.1 specification.

### TRAFFIC POLICING

Traffic policing is also called usage parameter control (UPC), and is performed at the input port of ATM switches to ensure that traffic generated by the user is within the negotiated contract. Should the contract be violated, the ATM switch of the



■ Figure 8. Protocol stack of an ATM hub and an ATM adapter card.



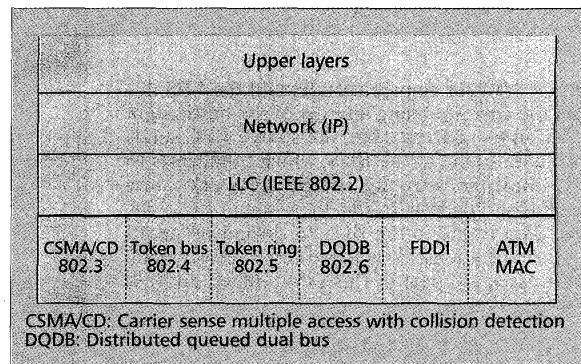
■ Figure 9. OSI and ATM protocol layer reference models.

ATM network service has the option of either discarding nonconforming ATM cells or tagging them as nonconforming by setting the cell loss priority (CLP) bit to 1 in the ATM header. The leaky bucket scheme is used to performed the policing function. For CBR traffic, a single leaky bucket is required to perform traffic policing, since CBR traffic uses a constant PCR parameter in its network contract. Traffic policing for VBR traffic utilizes two leaky buckets to monitor both the SCR over a defined period of time and the PCR used by the connection. If either of these parameters is violated, the ATM switch may discard the cell or mark it as nonconforming.

#### CONGESTION CONTROL

Congestion control is performed by the ATM network of switches to alleviate congestion inside the network. Since the network has allocated guaranteed resource (bandwidth and buffer) for CBR and VBR traffic, their QoS requirements can be satisfied. Since only the MCR is guaranteed for ABR, it may experience congestion when transmitting at a higher rate<sup>4</sup>; therefore, it is essential to inform the sources of ABR traffic when congestion occurs so that they may take appropriate action. A congestion control scheme based on feedback uses a closed-loop feedback control mechanism that allows the network to control the cell emission process at each source. Each virtual connection must have an independent control loop since each connection may follow a different path through the network. Two classes of feedback schemes have been proposed: *credit-based* and *rate-based*. The credit-based approach is the link-by-link window flow control scheme. Each link in the network independently runs the flow control mechanism [16]. Rate-based schemes use feedback information from the network to control the rate at which each source emits cells into the network on every virtual connection. Basically, the source sends a resource management (RM) cell periodically among data cells (or before a threshold time expires) to the destination, and the destination returns all the received RM cells to the source. The RM cells contain important information such as current cell rate (CCR), minimum cell rate (MCR), explicit rate (ER), and congestion indication (CI). The ABR traffic source always sends at a rate at least MCR, but never exceeding PCR, and adjusts its rate accordingly through the feedback information (ER and CI) from returning RM cells. Should the RM cell not return within a threshold time, the ABR source will decrease its rate accordingly. During congestion at a switching node, the switch uses at least one of the following methods to indicate congestion:

<sup>4</sup> Since no QoS parameters are specified for UBR, ATM cells from UBR traffic may be subject to discarding during network congestion.



■ Figure 10. LAN emulation (LE).

- The switching node sets the explicit forward congestion indication (EFCI) state in the ATM *data* cell headers.
- The switching node sets CI = 1 in the forward and/or backward RM cells.
- The switching node reduces the ER field of the forward and/or backward RM cells.

On the destination side, the host shall return all RM cells back to the source. Upon detection of EFCI = 1 in the ATM data cell that arrives just before an RM cell, the destination will mark the CI bit to 1 or reduce the ER level in the backward RM cell.

#### SUPPORTING TCP/IP OVER ATM NETWORKS

There are more than four million hosts on the Internet today, many running the TCP/IP protocol suite. Figure 9 shows the Open System Interconnection (OSI) and ATM protocol layer reference models. A successful introduction strategy for ATM-based LAN requires that existing services and protocols (i.e., TCP/IP) run transparently over ATM networks.

There are currently two schemes to support the TCP/IP protocol suite over ATM, namely *LAN Emulation* and *IP over ATM*.

#### LAN EMULATION

This defines a service interface for higher-layer protocols (logical link control — LLC — and IP). Under this scheme, a new ATM MAC sublayer is needed beneath the LLC sublayer (Fig. 10) and gives the appearance of a virtual shared medium such as IEEE 802.x LAN. As standardized by the ATM Forum [17], a LAN emulation service consists of the following three components:

- LAN emulation server (LES) — provides a facility for registration and resolving a MAC address into an ATM address
- Broadcast/unknown server (BUS) — forwards multicast/broadcast frames and delivers unicast frames for an unregistered or address-unresolvable LAN host
- LAN emulation configuration server (LECS) — locates the LES and obtains configuration information for each ATM segment

The three components of LE service can be implemented on a single physical entity or distributed on several physical entities. A device attached directly to an ATM switch (e.g., workstation, ATM hub, router) contains a LAN Emulation Client (LEC) (Fig. 11). The main function of the LEC is to communicate with remote LE service components (i.e., LES, BUS, and LECS). The LEC is identified by two addresses: a unique 6-byte MAC address and a 20-byte ATM address. For unicast traffic, the source LEC first sends an LE-ARP request to LES, asking for the ATM address of the destination LEC. With the LE-ARP reply from LES, the LEC establishes a direct VC connection to that ATM address, and the data

frame transfer commences. Should the LES not get the corresponding address of the destination LEC in its address table, the source LEC sends the data frame to the BUS, which broadcasts to all stations on the ATM network. For multicast/broadcast traffic, the source LEC sends the data frame to the BUS for broadcast to all attached stations. Only those LECs whose MAC addresses are part of this group MAC address take the data frame.

### IP OVER ATM

We may implement classical IP directly over ATM [18], in which case ATM is used as a direct replacement for the "wires" and local LAN segments connecting IP end-stations and routers operating in the "classical" LAN-based paradigm. Hosts of different IP subnets must communicate via an intermediate IP router, even though it may be possible to open a direct VC between the two IP end-stations over the ATM network. Obviously, this approach poses serious limitations. Currently, the Routing Over Larger Clouds (ROLC) working group of the Internet Engineering Task Force (IETF) is investigating the possibility of setting up direct connections across ATM [19]. With this aim, the ROLC working group is working on a new protocol named "Next Hop Resolution Protocol" (NHRP) which relies on the use of "super" servers. The ultimate goal of the NHRP protocol is to enable a host to bypass some or all of the routers between source and destination hosts by establishing a direct connection through the ATM switches.

### CONCLUSIONS AND FINAL THOUGHTS

In this article, we have covered key networking components and issues of an ATM-based enterprise network, including the ATM switch, host-network interface, traffic service classes and management functions, and software protocols for a smooth introduction of ATM into the desktop computing environment.

Early experience in broadband networking shows that data traffic (e.g., information browsing) will continue to be the dominant traffic volume for desktop applications. This reiterates the necessity of seamless integration of existing protocols (TCP/IP) in an ATM environment. As the ATM Forum rapidly puts ATM standards into print, we expect to see more network vendors enter the ATM marketplace. Also, to compete successfully with alternative networking products such as 100Base-T Ethernets, switched Ethernets, and fiber/copper distributed data interface (FDDI/CDDI), it is essential for ATM vendors to further cut the costs of ATM products to be more price-competitive. Currently, vendors are introducing the use of unshielded twisted pair (UTP) category 5 cable for 155 Mb/s. This will eliminate the relative expensive operation of electro-optical conversion at both the ATM adapter cards and ATM switch ports, which is expected to bring down the cost of ATM products significantly.

### ACKNOWLEDGMENTS

The first author would like to thank David A. Berkley, Steven Y. Gao and Frank C. Pirz of AT&T Bell Labs, Murray Hill, NJ for many insightful discussions on ATM LAN implementation.

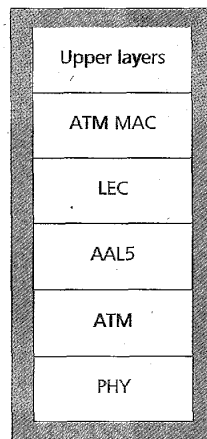


Figure 11. LAN Emulation Client (LEC) protocol layering.

### REFERENCES

- [1] T. G. Robertazzi, ed., *Performance Evaluation of High-Speed Switching Fabrics and Networks*, [New York: IEEE Press, 1993].
- [2] H. J. Chao, "Lecture Notes on Broadband Packet Switching," unpublished, Dec. 1994.
- [3] P. Newman, "ATM Technology for Corporate Networks," *IEEE Commun. Mag.*, Apr. 1992, pp. 90-101.
- [4] R. Roolamini et al., "Finding the Right ATM Switch for the Market," *IEEE Comp.*, Apr. 1994, pp. 17-28.
- [5] SynOptics Communications, Inc. "LattisCell," SynOptics Product Description, ID Number: PU470-1672US-A, Sept. 8, 1994.
- [6] FORE Systems, Inc. "White Paper: ATM Switching Architecture," Nov. 1993.
- [7] H. J. Chao, "A Recursive Modular Terabit/Second ATM Switch," *IEEE JSAC*, vol. 9, no. 8, Oct. 1991, pp. 1161-72.
- [8] M. De Prycker, *Asynchronous Transfer Mode: Solution to Broadband ISDN*, 2nd ed., [Ellis Horwood, 1993].
- [9] H. J. Chao and D. E. Smith, "A Shared-Memory Virtual Channel Queue for ATM Broadband Terminal Adapters," *Int'l. J. Digital and Analog Commun. Sys.*, vol. 5, 1992, pp. 29-37.
- [10] *IEEE JSAC*, Special Issue on High Speed Computer/Network Interfaces, Feb. 1993, vol. 11, no. 2.
- [11] P. A. Steenkiste, "A Systematic Approach to Host Interface Design for High-Speed Networks," *IEEE Comp.*, Mar. 1994, pp. 47-57.
- [12] FORE Systems, Inc., "200-Series ATM Adapter: Design and Architecture," 1995.
- [13] T. F. La Porta et al., "B-ISDN: A Technology Discontinuity," *IEEE Commun. Mag.*, Oct., 1994, pp. 84-97.
- [14] Bay Networks, Inc. "EtherCell Ethernet-to-ATM Switching," Dec. 19, 1994.
- [15] ATM Forum, "Traffic Management Specification - Version 4.0," atmf95-0013R7, June 1995.
- [16] H. T. Kung and R. Morris, "Credit-Based Flow Control for ATM Networks," *IEEE Network*, Mar./Apr. 1995, pp. 40-48.
- [17] ATM Forum, LAN Emulation SWG Drafting Group, "LAN Emulation Over ATM Specification - Version 1.0," Jan. 1995.
- [18] M. Laubach, "Classical IP and ARP over ATM," IETF RFC 1577, Apr. 1994.
- [19] N. Kavak, "Data Communication in ATM Networks," *IEEE Network*, May/June 1995, pp. 28-37.

### BIOGRAPHIES

YIWEI THOMAS HOU received the B.E. degree (Summa Cum Laude) from the City College of New York in 1991 and the M.S. degree from Columbia University in 1993, both in electrical engineering. Currently, he is a Ph.D. candidate in electrical engineering at Polytechnic University, Brooklyn, New York, under an NSF Graduate Research Traineeship. Mr. Hou spent the summers of 1994 and 1995 at AT&T Bell Laboratories, Murray Hill, New Jersey, working on ATM LAN internetworking. His research interests include modeling, analysis, optimization, and design in an ATM and TCP/IP internetworking environment.

LEANDROS TASSIULAS received the Diploma in electrical engineering from the Aristotelian University of Thessaloniki, Thessaloniki, Greece in 1987, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1989 and 1991, respectively. From September 1991 to June 1995 he was an assistant professor in the Department of Electrical Engineering, Polytechnic University. Since July 1995 he has been an assistant professor in the Department of Electrical Engineering at University of Maryland, College Park.

H. JONATHAN CHAO received the B.S.E.E. and M.S.E.E. degrees from National Chiao Tung University, Taiwan, in 1977 and 1980, respectively, and the Ph.D. degree in electrical engineering from Ohio State University in 1985. From 1977 to 1981, he worked at Taiwan Telecommunication Laboratories, where he was engaged in the development of a digital switching system. From 1985 to 1991, he was a member of technical staff at Bellcore. Currently, he is an associate professor of electrical engineering at Polytechnic University, which he joined in January 1992. He holds 14 patents and has published more than 50 journal and conference papers in the above areas.